Optimal individualized decision rules using instrumental variable methods

Hongxiang Qiu[1]

Joint work with:
Marco Carone[1], Ekaterina Sadikova[2], Maria Petukhova[2], Ronald C. Kessler[2], Alex Luedtke[3]

[1]: Dept. of Biostatistics, University of Washington
[2]: Dept. of Health Care Policy, Harvard Medical School
[3]: Dept. of Statistics, University of Washington

My research

My web page: https://Qiu-Hongxiang-David.github.io

Qiu H, Carone M, Sadikova E, Petukhova M, Kessler R, Luedtke A (2020). Optimal individualized decision rules using instrumental variable methods. *Journal of the American Statistical Association* (with discussion), 1-18.

Qiu H, Luedtke A, Carone M (2020). Universal sieve-based strategies for efficient estimation using machine learning tools. *arXiv.* (accepted by *Bernoulli*)

Qiu H, Luedtke A (2021+). Leveraging vague prior information in general models via iteratively constructed Gamma-minimax estimators. *arXiv.*

Bobb J, **Qiu H**, Matthews, A, McCormack J, Bradley K (2020). Addressing identification bias in the design and analysis of cluster-randomized pragmatic trials: a case study. *Trials* 21(1), 289.

Motivation

- Emerging area of individualized treatment rules (ITR).
- Previous methods assume no unmeasured confounding (e.g., Murphy, 2003; Robins, 2004; Zhao et al., 2012; Chakraborty and Moodie, 2013; Luedtke and van der Laan, 2016b).
  What if there is unmeasured confounding?

Motivation

- Emerging area of individualized treatment rules (ITR).
- Previous methods assume no unmeasured confounding (e.g., Murphy, 2003; Robins, 2004; Zhao et al., 2012; Chakraborty and Moodie, 2013; Luedtke and van der Laan, 2016b).
  What if there is unmeasured confounding?
- Instrumental variable (IV): another approach to identifying causal effects.
  Can we use an IV to estimate an optimal ITR?

Motivation

- Emerging area of individualized treatment rules (ITR).
- Previous methods assume no unmeasured confounding (e.g., Murphy, 2003; Robins, 2004; Zhao et al., 2012; Chakraborty and Moodie, 2013; Luedtke and van der Laan, 2016b).
  What if there is unmeasured confounding?
- Instrumental variable (IV): another approach to identifying causal effects.
  Can we use an IV to estimate an optimal ITR?
- Example:
  - IV: randomized treatment assignment
  - Treatment: actual treatment status
- Especially interested in settings with a treatment resource constraint (Luedtke and van der Laan, 2016a).

Motivation

- At times, direct intervention on treatment may be impossible or expensive.
- Individualized encouragement rule (IER): intervention on IV.
- Evaluate optimal rules: average benefit under optimal rule (compared to a reference rule).

Table of Contents

1. Setup and causal estimands
2. Identifying conditions and results
3. Estimation and inference under a locally nonparametric model
4. Results

Problem setup

- Observe iid $O = (W, Z, A, Y) \sim P_0$:
  - $W \in \mathcal{W}$: baseline covariates.
  - $Z \in \{0, 1\}$: binary IV.
  - $A \in \{0, 1\}$: binary treatment status (treatment vs control).
  - $Y \in \mathbb{R}$: outcome of interest (larger values are preferable).
- (Stochastic) individualized rule: $d : \mathcal{W} \to [0, 1]$ (prob of treatment)
- Counterfactuals:
  - $A(z)$: potential treatment status corresponding to $Z = z$
  - $Y(z, a)$: potential outcome corresponding to $(Z, A) = (z, a)$
- Given treatment resource constraint: $P_0(\text{receiving treatment}) \leq \kappa$.

IV conditions

- Relevance: $|P_0(A=1 \mid Z=1,W) - P_0(A=1 \mid Z=0,W)| > \delta^A$
- Exclusion restriction: $Y(0,a) = Y(1,a) =: Y(a)$
- Independence: $Z \perp\!\!\!\perp U \mid W$
- Z: IV/encouragement.
- U: unobserved confounder.

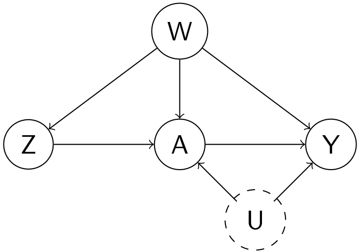

Case I: intervention on treatment

For an ITR $t : \mathcal{W} \to [0, 1]$, $Y(t) :=$ counterfactual outcome under t.

The optimal ITR t_0 solves

$$\textbf{maximize} \;\; \mathbb{E}[Y(t)] \;\; \textbf{subject to} \;\; \mathbb{E}[t(W)] \leq \kappa \; .$$

The impact of implementing the optimal ITR can be measured via its average treatment effect (ATE) relative to a given reference ITR t_r:

$$\mathbb{E}[Y(t_0) - Y(t_r)]$$

Case II: intervention on encouragement

For an IER $e : \mathcal{W} \to [0,1]$, $A(e) :=$ counterfactual treatment under e.

The optimal IER e_0 solves

$$\textbf{maximize} \;\; \mathbb{E}[Y(A(e))] \;\; \textbf{subject to} \;\; \mathbb{E}[A(e)] \leq \kappa \; .$$

We intervene on the IV but the constraint is on the (stochastic!) treatment.

The impact of implementing the optimal IER can be measured via its average encouragement effect (AEE) relative to a given reference IER e_r:

$$\mathbb{E}[Y(A(e_0)) - Y(A(e_r))]$$

Table of Contents

1 Setup and causal estimands

2 Identifying conditions and results

3 Estimation and inference under a locally nonparametric model

4 Results

Additional notations

- IV $\rightarrow$ outcome effect (conditional average encouragement effect):

$$\Delta_0^Y(w) = \mathrm{E}_0[Y|Z=1, W=w] - \mathrm{E}_0[Y|Z=0, W=w]$$

- Wald estimand (conditional average treatment effect):

$$\Delta_0(w) = \frac{\mathrm{E}_0[Y|Z=1, W=w] - \mathrm{E}_0[Y|Z=0, W=w]}{\mathrm{E}_0[A|Z=1, W=w] - \mathrm{E}_0[A|Z=0, W=w]}$$

- Proportion treated among encouraged:

$$\mathrm{E}_0[A|Z=1, W=w]$$

Case I: intervention on treatment

Key identifying conditions (slightly relaxed version of Wang and Tchetgen Tchetgen (2018)):

- $Y(a) \perp\!\!\!\perp (A, Z) \mid (W, U)$.
- One of the following sets of conditions holds:

(1) (a) (Uncorrelated IV) $\mathrm{Cov}(Y(0), Z \mid W) = 0$

(b) (No unmeasured treatment-outcome effect modification)

$$\mathbb{E}[Y(1) - Y(0) \mid W, U] = \mathbb{E}[Y(1) - Y(0) \mid W]$$

(2) (a) (Independent IV) $Z \perp\!\!\!\perp U \mid W$

(b) (Independent compliance)

$$\mathbb{E}[A(Z) \mid Z = 1, W, U] - \mathbb{E}[A(Z) \mid Z = 0, W, U]$$
$$= \mathrm{E}_0[A \mid Z = 1, W] - \mathrm{E}_0[A \mid Z = 0, W]$$

Case I: average treatment effect

ATE can be written as a summary of P_0.

Theorem (Identification of ATE)

- $\mathbb{E}[Y(1) - Y(0) \mid W] = \Delta_0(W)$
- $\mathbb{E}[Y(t) - Y(t_r)] = \mathrm{E}_0[\{t(W) - t_r(W)\}\,\Delta_0(W)]$ for any ITR t

Case I: optimal ITR

What does an optimal ITR look like?

Case I: optimal ITR

What does an optimal ITR look like?

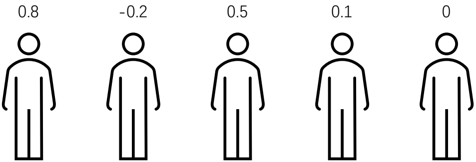

Case I: optimal ITR

1. Sort subgroups according to $\Delta_0(W)$ (from high to low).
2. Assign treatment to those with highest (and positive) conditional ATE $\Delta_0(W)$ until treatment runs out.

$$t_0(w) = I\left\{\Delta_0(w) > \tau_0^T\right\}$$

Case II: intervention on encouragement

Key identifying condition: $Z \perp\!\!\!\perp (A(z), Y(a)) \mid W$

AEE can be written as a summary of P_0.

Theorem (Identification of AEE)

- $\mathbb{E}[Y(A(1)) - Y(A(0)) \mid W] = \Delta_0^Y(W)$
- $\mathbb{E}\left[Y(A(e)) - Y(A(e_r))\right] = \mathrm{E}_0\left[\{e(W) - e_r(W)\}\, \Delta_0^Y(W)\right]$ for any IER e

Case II: optimal IER

What does an optimal IER look like?

Case II: optimal IER

What does an optimal IER look like?

View as a fractional knapsack problem:

- Subgroup with same W: item
- Conditional AEE: $\Delta_0^Y(W)$ = value
- Proportion treated among encouraged: $E_0[A|Z=1,W]$ = weight
- κ = total weight capacity
- $\xi_0(W) := \Delta_0^Y(W)/\bar{\mu}_0^A(W)$ = unit value

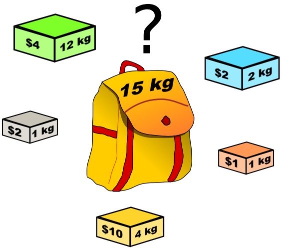

Case II: optimal IER

1. Sort subgroups according to $\xi_0(W)$ from high to low
2. Assign encouragement to treatment to those with highest (and positive) $\xi_0(W)$ until treatment runs out

$$e_0(w) = I\left(\xi_0(w) > \tau_0^E\right)$$

Table of Contents

1. Setup and causal estimands

2. Identifying conditions and results

3. Estimation and inference under a locally nonparametric model

4. Results

Overview of targeted minimum-loss based estimation

- Goal: estimate $\Phi(P_0)$
- Seemingly natural approach:
  - Estimate P_0 with $\tilde{P}_n$
  - Plug in: estimate $\Phi(P_0)$ with $\Phi(\tilde{P}_n)$
- Problem: typically inefficient and not asymptotically normal.

Overview of targeted minimum-loss based estimation

- In first order, bias of $\Phi(\tilde{P}_n) \approx -\frac{1}{n}\sum_{i=1}^{n} G(\tilde{P}_n)(O_i)$
  - $G(P)$: canonical gradient at P. Also called efficient influence function.
- One-step correction: $\Phi(\tilde{P}_n) + \frac{1}{n}\sum_{i=1}^{n} G(\tilde{P}_n)(O_i)$
- TMLE algorithm:
  1. Targeting: fluctuate $\tilde{P}_n$ and find $\hat{P}_n$ such that the approximated bias is 0:

     $$\frac{1}{n}\sum_{i=1}^{n} G(\hat{P}_n)(O_i) = 0$$

     (Often by running a regression with clever covariates)
  2. Plug in: estimator given by $\Phi(\hat{P}_n)$

TMLE based on pseudo-gradient

- For our problem, it is difficult to solve $\frac{1}{n}\sum_{i=1}^{n} G(\hat{P}_n)(O_i) = 0$
  - May need to fit regressions iteratively
  - May lead to computational and statistical challenges
- We use TMLE based on pseudo-gradients to overcome these challenges
  - Pseudo-gradient: $G_{\tilde{P}_n}(P) \approx G(P)$ when $P \approx \tilde{P}_n$
  - Constructed by replacing some parts of P appearing in $G(P)$ with those of $\tilde{P}_n$
  - Fluctuate $\tilde{P}_n$ and find $\hat{P}_n$ such that $\frac{1}{n}\sum_{i=1}^{n} G_{\tilde{P}_n}(\hat{P}_n)(O_i) = 0$.
  - No need for iteratively fitting regressions.

Case I: intervention on treatment

Steps of proposed procedure:

1. Flexibly estimate relevant regression functions.
2. Estimate t_0 with sample analogue t_n.
3. Targeted estimation of $\psi_0^T := \mathrm{E}_0\left[\{t_0(W) - t_r(W)\}\,\Delta_0(W)\right]$

Case I: intervention on treatment

Key conditions:

- Non-exceptional law: $P_0\left\{\Delta_0(W) = \tau_0^T\right\} = 0$ (Robins, 2004)
- Function estimators converge sufficiently fast

Case I: intervention on treatment

Theorem (Asymptotic normality of ψ_n^T)

$$\psi_n^T - \psi_0^T = \frac{1}{n}\sum_{i=1}^{n} D^T(P_0)(O_i) + o_p(n^{-1/2})$$

where the influence function takes the form

$$D^T(P_0)(o) = \frac{t_0(w) - t_r(w)}{\Delta_0^A(w)[z + \mu_0^Z(w) - 1]}\left\{y - \mu_0^Y(z,w) - \Delta_0(w)[a - \mu_0^A(z,w)]\right\} + \left\{[t_0(w) - t_r(w)]\Delta_0(w) - \psi_0^T\right\} - \tau_0^T[t_0(w) - \kappa]$$

Case II: intervention on encouragement

Steps of proposed procedure:

1. Flexibly estimate relevant regression functions.
2. Estimate e_0 with sample analogue e_n.
   - If the estimated threshold>0, find a refined estimator k_n of κ
   - e_n is more likely to exactly respect the constraint κ
   - Needed for solving estimating equation.
3. Targeted estimation of $\psi_0^E := \mathrm{E}_0\left[\{e_0(W) - e_r(W)\}\, \Delta_0^Y(W)\right]$

Case II: intervention on encouragement

Key conditions:

- Non-exceptional law: $P_0(\xi_0(W) = \tau_0^E) = 0$ (Robins, 2004).
- Function estimators converge sufficiently fast.

Case II: intervention on encouragement

Theorem (Asymptotic normality of ψ_n^E)

$$\psi_n^E - \psi_0^E = \frac{1}{n}\sum_{i=1}^{n} D^E(P_0)(O_i) + o_p(n^{-1/2})$$

where the influence function takes the form

$$D^E(P_0)(o) = \frac{e_0(w) - e_r(w)}{z + \mu_0^A(w) - 1}[y - \mu_0^Y(z, w)] + [e_0(w) - e_r(w)]\Delta_0^Y(w) - \psi_0^E$$
$$- \tau_0^E \left[e_0(w) \left\{ \frac{z}{\mu_0^Z(w)}[a - \mu_0^A(1, w)] + \mu_0^A(1, w) \right\} - \kappa \right]$$

Table of Contents

1. Setup and causal estimands
2. Identifying conditions and results
3. Estimation and inference under a locally nonparametric model
4. Results

Simulation

Estimands:

- Resource constrained AEE: AEE of e_0 with $\kappa = 0.25$
- ATE: ATE of t_0 without resource constraint ($\kappa = 1$)
- Resource constrained ATE: ATE of t_0 with $\kappa = 0.25$

Data-generating mechanism has strong unmeasured treatment-outcome confounding.

Use sample splitting when estimating the optimal rule

- avoid a main source of finite-sample positive bias
- possible finite-sample negative bias
- valid 97.5% confidence lower bound, even under poorly estimated optimal rule

Simulation: AEE with resource constraints

Performance measure	Sample size	IV
95% Wald CI coverage	200	78%
	500	76%
	1000	74%
	2000	78%
97.5% confidence lower bound coverage	200	96%
	500	96%
	1000	96%
	2000	98%

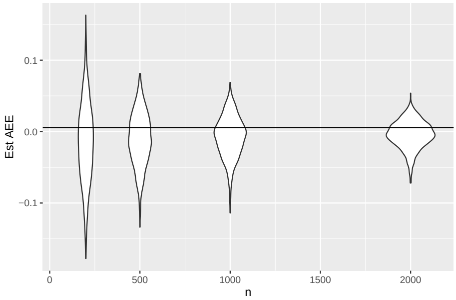

Simulation: ATE

Performance measure	Sample size	IV	Confounder
95% Wald CI coverage	200	97%	3%
	500	95%	$<1\%$
	1000	93%	$<1\%$
	2000	92%	$<1\%$
97.5% confidence lower bound coverage	200	$>99\%$	3%
	500	$>99\%$	$<1\%$
	1000	$>99\%$	$<1\%$
	2000	$>99\%$	$<1\%$

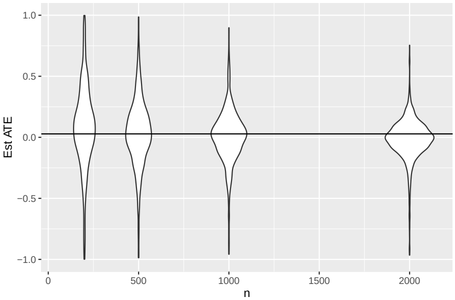

Simulation: ATE with resource constraints

Performance measure	Sample size	IV	Confounder
95% Wald CI coverage	200	96%	83%
	500	95%	84%
	1000	93%	87%
	2000	94%	88%
97.5% confidence lower bound coverage	200	$>$ 99%	97%
	500	$>$ 99%	97%
	1000	$>$ 99%	96%
	2000	$>$ 99%	92%

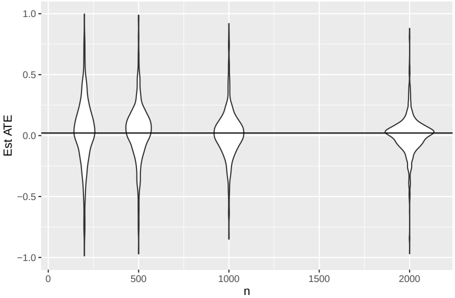

US army data

- W: 17 baseline binary covariates of soldiers in the month before the start of deployment.
- Z: indicator of initial assignment to a unit with low ($<$50%) probability of deployment
- A: indicator of non-deployment
- Y: indicator of no suicide-related outcomes*
- 55,272 soldiers, 38,404 in low POD units ($Z = 1$), 2.3% with $Y = 0$

*indicator of the soldier having neither died by suicide, made a nonfatal suicide attempt, been hospitalized for a psychiatric disorder, or had medically-reported suicidal ideation in the first 24 months post-deployment

US army data

Table: Estimates and approximate 95% CIs for the ATE and AEE of optimal ITRs and IERs at different constraint levels on the proportion of soldiers deployed in the unit.

min deployment	ATE
50%	0.021% (-0.211%, 0.254%)
75%	-0.056% (-0.271%, 0.160%)
90%	0.053% (-0.132%, 0.238%)

min deployment	AEE
50%	0.045% (-0.088%, 0.178%)
75%	0.087% (-0.052%, 0.227%)
90%	-0.042% (-0.162%, 0.078%)

Discussion

- Estimators of optimal individualized treatment/encouragement rule using an IV under treatment resource constraints.
- Inference on average causal effects of the optimal rule under a locally nonparametric model.
- Cui and Tchetgen Tchetgen (2020) studied a similar problem.
  - Both are discussion papers in the same issue of *JASA*.
  - Cui needed not consider intervention on encouragement or resource constraints.
  - Cui's identifying conditions for optimal ITR (not ATE) are weaker.
- Longitudinal setting: dynamic treatment?

References

Chakraborty, B. and E. E. Moodie (2013). *Statistical Methods for Dynamic Treatment Regimes*. Statistics for Biology and Health. New York, NY: Springer New York.

Cui, Y. and E. Tchetgen Tchetgen (2020). A Semiparametric Instrumental Variable Approach to Optimal Treatment Regimes Under Endogeneity. *Journal of the American Statistical Association*.

Luedtke, A. R. and M. J. van der Laan (2016a). Optimal Individualized Treatments in Resource-Limited Settings. *International Journal of Biostatistics 12*(1), 283–303.

Luedtke, A. R. and M. J. van der Laan (2016b). Statistical inference for the mean outcome under a possibly non-unique optimal treatment strategy. *Annals of Statistics 44*(2), 713–742.

Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 65*(2), 331–355.

Robins, J. M. (2004). Optimal Structural Nested Models for Optimal Sequential Decisions. pp. 189–326. Springer, New York, NY.

Wang, L. and E. Tchetgen Tchetgen (2018). Bounded, efficient and multiply robust estimation of average treatment effects using instrumental variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 80*(3), 531–550.

Zhao, Y., D. Zeng, A. J. Rush, and M. R. Kosorok (2012). Estimating Individualized Treatment Rules Using Outcome Weighted Learning. *Journal of the American Statistical Association 107*(499), 1106–1118.

Thank you!

Questions?

Additional notations

For any distribution P on observed data*, define

- $\mu_P^Z : w \mapsto \mathrm{E}_P(Z \mid W = w)$ *(IV propensity)*
- $\mu_P^A : (z, w) \mapsto \mathrm{E}_P\left(A \mid Z = z, W = w\right)$ *(Treatment regression)*
- $\mu_P^Y : (z, w) \mapsto \mathrm{E}_P\left(Y \mid Z = z, W = w\right)$ *(Outcome regression)*
- $\Delta_P^A : w \mapsto \mu_P^A(1, w) - \mu_P^A(0, w)$ *(IV $\to$ treatment effect)*
- $\Delta_P^Y : w \mapsto \mu_P^Y(1, w) - \mu_P^Y(0, w)$ *(IV $\to$ outcome effect)*
- $\Delta_P : w \mapsto \Delta_P^Y(w)/\Delta_P^A(w)$ *(Wald estimand)*
- $\bar{\mu}_P^A : w \mapsto \mu_P^A(1, w)$ *(Proportion of following encouragement 1)*

*For P_0, we use 0 instead of P_0 in the subscript.

Formal optimal ITR

Theorem (Identifying optimal ITR)

With $\eta_0^T := \inf\{\eta : P_0(\Delta_0(W) > \eta) \leq \kappa\}$ *and the threshold* $\tau_0^T := \max\{\eta_0^T, 0\}$, *the optimal ITR is*

$$t_0(w) := \begin{cases} \frac{\kappa - P_0\{\Delta_0(W) > \tau_0^T\}}{P_0\{\Delta_0(W) = \tau_0^T\}} & : \Delta_0(w) = \tau_0^T > 0, \\ & \quad P_0\{\Delta_0(W) = \tau_0^T\} > 0 \\ I\{\Delta_0(w) > \tau_0^T\} & : \text{otherwise.} \end{cases}$$

Formal optimal IER

Theorem (Identifying optimal IER)

With $\eta_0^E := \inf\left\{\eta : \mathrm{E}_0\left[I(\xi_0(W) > \eta)\bar{\mu}_0^A(W)\right] \leq \kappa\right\}$ *and* $\tau_0^E := \max\left\{\eta_0^E, 0\right\}$, *the optimal IER is*

$$e_0(w) = \begin{cases} \frac{\kappa - \mathrm{E}_0\left[I(\xi_0(W) > \tau_0^E)\bar{\mu}_0^A(W)\right]}{\mathrm{E}_0\left[I(\xi_0(W) = \tau_0^E)\bar{\mu}_0^A(W)\right]} & : \xi_0(w) = \tau_0^E > 0, \\ & \quad \mathrm{E}_0\left[I(\xi_0(W) = \tau_0^E)\bar{\mu}_0^A(W)\right] > 0 \\ I\left(\xi_0(w) > \tau_0^E\right) & : \textit{otherwise.} \end{cases}$$

Case I: intervention on treatment

Proposed procedure:

1. Estimate relevant regression functions via machine learning and compute Δ_n.
2. Estimate t_0 with the sample analogue t_n using Δ_n.
3. Targeted estimation of $\psi_0^T := \mathrm{E}_0\left[\{t_0(W) - t_r(W)\}\,\Delta_0(W)\right]$:
   - (a) Obtain a targeted estimate $\hat{\mu}_n^Y$ of μ_0^Y by running an ordinary least-squares linear regression with outcome Y, covariate $h(Z, W) := \frac{t_n(W) - t_r(W)}{[Z + \mu_n^Z(W) - 1]\Delta_n^A(W)}$, offset $\mu_n^Y(Z, W)$ and no intercept;
   - (b) Obtain a targeted estimate $\hat{\mu}_n^A$ of μ_0^A by running a logistic regression with outcome A, covariate $h(Z, W)\Delta_n(W)$, offset logit $\mu_n^A(Z, W)$ and no intercept;
   - (c) Estimate ψ_0^T with $\psi_n^T := \frac{1}{n}\sum_{i=1}^n \{t_n(W_i) - t_r(W_i)\}\hat{\Delta}_n(W_i)$.

Case II: intervention on encouragement I

Proposed procedure:

1. Estimate relevant regression functions via machine learning and compute $\Delta_n^Y, \bar{\mu}_n^A$.
2. Estimate e_0 using the sample analogue with a refined estimator k_n of κ:
   (a) For any $k \in [0,1]$, let $\eta_n^E(k)$, $\tau_n^E(k)$ and $d_{n,k}$ be the sample analogues for the optimal IER when the constraint is k (except that $d_{n,k}$ uses threshold $\eta_n^E(k)$ rather than $\tau_n^E(k)$).
   (b) If $\tau_n^E(\kappa) > 0$ and there is a solution to
   $$\frac{1}{n}\sum_{i=1}^{n} d_{n,k}(W_i)\left\{\bar{\mu}_n^A(W_i) + \frac{I(Z_i=1)}{\mu_n^Z(W_i)}\left[I(A_i=1) - \mu_n^A(1,W_i)\right]\right\} = \kappa,$$
   set k_n to the solution; otherwise, set $k_n = \kappa$.
   (c) Estimate e_0 with its sample analogue e_n, except that the treatment resource constraint is represented by k_n.

Case II: intervention on encouragement II

3. Targeted estimation of $\psi_0^E := \mathrm{E}_0\left[\{e_0(W) - e_r(W)\}\,\Delta_0^Y(W)\right]$:
   (a) Obtain a targeted estimate $\hat{\mu}_n^Y$ of μ_0^Y by running an ordinary least-squares linear regression with outcome Y, covariate $[e_n(W) - e_r(W)]/[Z + \mu_n^Z(W) - 1]$, offset $\mu_n^Y(Z, W)$ and no intercept, and taking $\hat{\mu}_n^Y$ to be the fitted mean function.
   (b) Estimate ψ_0^E with $\psi_n^E := \frac{1}{n}\sum_{i=1}^n \{e_n(W_i) - e_r(W_i)\}\,\hat{\Delta}_n^Y(W_i)$.

Simulation setting

$$W := (W_1, W_2, W_3), W_1 \sim \mathrm{Unif}(-1,1), W_2 \sim \mathrm{Bern}(0.5), W_3 \sim \mathrm{N}(0,1),$$
$$U \sim \mathrm{Bern}(0.5),$$
$$Z \sim \mathrm{Bern}(\mathrm{expit}\{2.5W_1 + 0.5W_2W_3\}),$$
$$A \sim \mathrm{Bern}(0.6\,\mathrm{expit}\{2Z + W_1 - W_2 + 0.7W_3\} + 0.2 + 0.4(U - 0.5)),$$
$$Y \sim \mathrm{Bern}(\mathrm{expit}\{AW_1 + 0.2W_2 - 0.5W_3 + 4(U - 0.5)\}).$$

Machine learning: Super Learner with library including

- logistic regression
- generalized additive model with logit link
- gradient boosting